ESTABLISHING THE HUMAN FIREWALL: REDUCING AN INDIVIDUAL'S VULNERABILITY TO SOCIAL ENGINEERING ATTACKS

THESIS

Jamison W. Scheeres

AFIT/GIR/ENG/08-04

DEPARTMENT OF THE AIR FORCE
AIR UNIVERSITY

AIR FORCE INSTITUTE OF TECHNOLOGY

Wright-Patterson Air Force Base, Ohio

The views expressed in this thesis are those of the author and do not reflect the official policy or position of the United States Air Force, Department of Defense, or the U.S. Government.

AFIT/GIR/ENG/08-04

ESTABLISHING THE HUMAN FIREWALL: REDUCING AN INDIVIDUAL'S VULNERABILITY TO SOCIAL ENGINEERING

THESIS

Presented to the Faculty

Department of Electrical and Computer Engineering

Graduate School of Engineering and Management

Air Force Institute of Technology

Air University

Air Education and Training Command

In Partial Fulfillment of the Requirements for the

Degree of Master of Science in Systems Engineering

Jamison W. Scheeres, BS

March 2008

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED

AFIT/GIR/ENG/08-04

ESTABLISHING THE HUMAN FIREWALL: REDUCING AN INDIVIDUAL'S VULNERABILITY TO SOCIAL ENGINEERING ATTACKS

Jamison W. Scheeres, BS

Approved:

//signed// 5 Mar 08
Robert F. Mills, PhD (Chairman) Date

//signed// 5 Mar 08
Michael A. Grimalia, PhD (Member) Date

//signed// 5 Mar 08
Michael T. Rehg, PhD (Member) Date

AFIT/GIR/ENG/08-04

Abstract

Hackers frequently use social engineering attacks to gain a foothold into a target network. This type of attack is a tremendous challenge to defend against, as the weakness lies in the human users, not in the technology. Thus far, methods for dealing with this threat have included establishing better security policies and educating users on the threat that exists. Existing techniques aren't working as evidenced by the fact that auditing agencies consider it a given that will be able to gain access via social engineering. The purpose of this research is to propose a better method of reducing an individual's vulnerability to social engineering attacks.

The first part of the research formally establishes the connection between social engineering attacks and what social psychologists call, "illegitimate persuasion" using content analysis to show that both target the same common psychological triggers. This comparison is done by comparing specific examples in different mediums. Once this connection is established, this research also proposes how to apply specific techniques that have been shown to be effective in increasing an individual's resistance to persuasion. Specifically, it uses techniques that have been shown to be effective in bolstering an individual's resistance illegitimate persuasion. The research culminates in the proposal of a template for a training program using these new techniques that can be incorporated into existing training classes that should significantly reduce an individual's vulnerability to social engineering attacks.

To Emily

Acknowledgements

I would first like to thank my faculty advisor, Dr. Robert Mills for his expertise and guidance throughout this research effort. He provided invaluable dialogue, support, and alternatives when I thought I had hit some roadblocks. I would also like to thank my committee members, Dr. Michael R. Grimalia – for some of the ideas on persuasion and social engineering he lent to this project and for introducing me to my sponsor – and Dr. Michael T. Rehg – for some of his ideas on training and learning theory. Additionally, I would like to thank my sponsor, LT Julie Ann Janson from AFRL for believing in my work and sponsoring it; I hope it lives up to your expectations. I would also like to submit a very special thank you to Dr. Brad Sagarin for taking some time out to talk with me and discuss some of his perspectives on this work, your insights were invaluable and I hope we can continue our discussions in the future. I would also like to thank Major Keile Weigle of the 39 IOS and Major Chris Evans of the 57 IAS for the unique insights they and their teams were able to give me on social engineering. Finally, I would like to thank my family: My Mom for the insights she gave and discussions we were able to have during the process of research; my Dad for his love and support; and my wife for her absolute and unconditional love and support through this very challenging task; she helped make it the most rewarding. Most importantly, I would like to thank God for guiding me through this tumultuous last year. Truly all things are possible through Him.

Jamison W. Scheeres

Table of Contents

Page

List of Figures

Figure Page

ESTABLISHING THE HUMAN FIREWALL: REDUCING AN INDIVIDUAL'S VULNERABILITY TO SOCIAL ENGINEERING

I. Introduction

1.1 Overview

Human users are the weakest link in any given information security system (Mitnick and Simon, 2002); they can be deceived, tricked, and manipulated into circumventing the security protocols in place. As technology progresses, it is this human weakness that remains the most challenging aspect of system security. Mitnick describes the techniques used to exploit this human vulnerability to bypass security systems in order to gather information as "social engineering" (Mitnick and Simon, 2002). Human users create holes in even the most expensive and well thought-out hardware and software protection system to such a point that, when conducting an audit of a company's information security policies, some security consultants consider it a given that they would be able to gain access to guarded information via social engineering (Gragg, 2003). Each new generation of technology promises a new and improved level of security, but as Peltier points out, the 1970s promised access control packages to make systems secure, the 1980s promised that effective anti-virus software would make systems secure, the 1990s promised that firewalls would make systems secure, and here in the 21st century, the promise is that intrusion detection systems and public key infrastructure are what's going to make our systems secure (Peltier, 2006). All may prove ineffective as attack technologies and techniques continue to evolve and adapt to new technology, but one method of attacking a network has remained constant, the attack on the users. As security technology improves its ability to keep out those who would infiltrate the network, perpetrators likely will simply shift their focus to the human user in order to gain access. Ironically, this human aspect of security has in the past been largely ignored or downplayed as being too difficult a problem to solve. The purpose of this thesis is to develop a means of framing and solving the social engineering problem.

1.2 Recognizing the Threat

Many organizations, including the Air Force, use vulnerability assessment teams that use all the tricks and tactics that a perpetrator would to conduct internal audits on their defenses. These so-called “red teams” almost always use social engineering techniques to create an initial point of entry into a network and then use that entry point as a jumping off point in order to work their way further in the organization’s physical and network resources. Once inside, the red teams can extract information or plant programs that create backdoors into the network for later use. Anecdotal evidence from interviews with past and present members of these red teams has indicated that these teams gain access to a network 100 percent of the time via social engineering techniques—social engineering *always* works! The threat of social engineering cannot be overstated, as it is an insidious, difficult-to-detect method of gaining access to information or resources that otherwise would be considered well-protected. An attack may consist of gaining seemingly innocuous information, combining it with seemingly innocuous information gathered from another source and suddenly some critical bit of information is revealed than can be used to exploit even more information from an unsuspecting user. This linkage can be seen in Figure 1, an illustration from a 1995 article on social engineering by Winkler and Dealy (1995). What this image helps to illustrate is that a determined social engineer is rarely after a single piece of information, but rather gathers what she can, when she can and then combines the collected pieces of information into a potential attack vector.

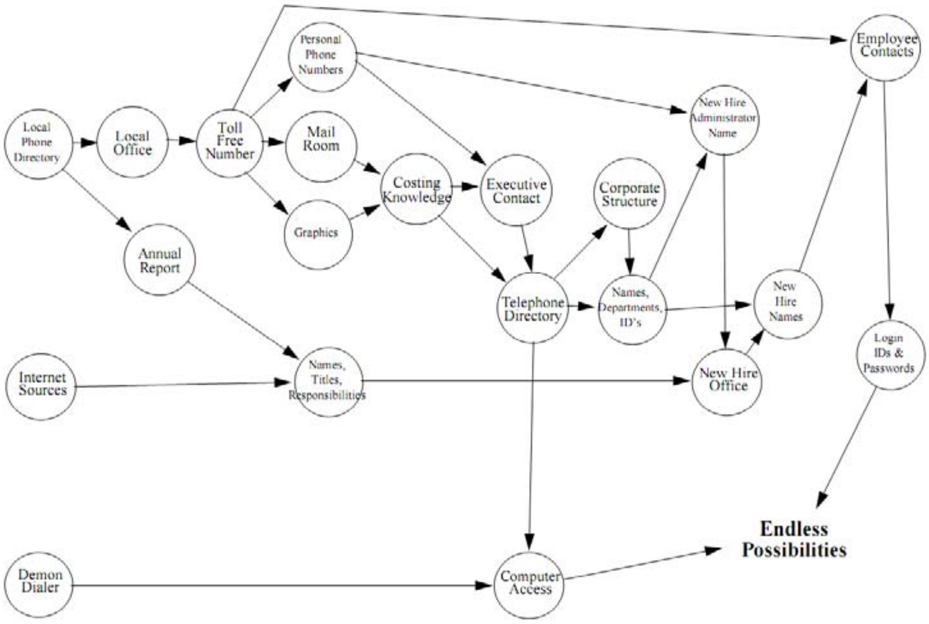

FIGURE 1: ANATOMY OF AN ATTACK (WINKLER AND DEALY, 1995)

1.3 Psychology and Social Engineering

The exact tactics and techniques used by social engineers are constantly evolving, but the common thread among the many techniques of social engineering is that they target certain "psychological triggers" to accomplish their objectives (Gragg, 2003, Rusch, 1999). Trade literature and general practitioner attitudes highlight the importance of user education and strong security policies as primary methods for defense against social engineering attacks (Gragg, 2003, Granger, 2002, Mitnick and Simon, 2002), but the continuing vulnerability of organizations despite the implementation of these policies indicate that something further is needed. Additional methods to combat social engineering are clearly needed, and need to be fielded as soon as possible. One area of study that may help in uncovering a better method of combating social engineering is the field of psychology, specifically social psychology.

Gragg (2003) has described several psychological triggers or vulnerabilities that social engineers exploit to create an emotional state in an individual that leaves them willing to grant access to information that would otherwise be protected. Inherent in the use of these triggers is deception. In almost every social engineering attack, the perpetrator wants the victim to believe the perpetrator is someone they are not. Similarly, a form of persuasion that has been termed "illegitimate persuasion" by Dr. Brad Sagarin and his colleagues, is when a message seeks "to convince through the use of deceptive tactics" such as having an actor dressed as a stock broker who recommends a particular brokerage firm, when in fact, the actor has no specific expertise to make such a recommendation (Sagarin et al., 2002). This concept dovetails well with the definition of social engineering. A key difference is that with illegitimate persuasion, the objective typically is to convince an individual to buy something, whereas with social engineering, the desired objective is convincing an individual to reveal sensitive information or bypass established security controls.

1.4 Resistance to Persuasion & Social Engineering

Persuasion is fundamental to human interaction, and social psychologists have long been interested in understanding which persuasive tactics are effective and why. Considerable research has looked at the forms of persuasion (Cialdini, 1993, Cialdini, 2001), while in contrast, relatively few have examined persuasion defenses. The forms of defense against persuasion that have been defined are based largely on the work of McGuire, who proposed an inoculation defense (McGuire, 1964), as well as research done by Petty and Caciappo on forewarning (Petty and Caciappo, 1977). Of these two persuasion defense theories, the inoculation theory appears to have possible applications to building a defense against social engineering. McGuire's inoculation theory suggests that individuals become resistant to social influences when they are exposed to less severe threats, and mirrors the medical practice of exposing a patient to a reduced concentration of a disease to build up antibodies to counter that disease. Moreover, an extension of McGuire's inoculation theory has been effective against illegitimate persuasion in an experimental setting (Sagarin et al., 2002) which showed that dispelling any illusions of invulnerability that an individual has is necessary to increasing an individual's resistance.

1.5 Academic Challenges of Studying Social Engineering

While trade literature on social engineering is plentiful, the topic is a difficult problem to tackle in an academic setting. Techniques of social engineering are observed only after an attack has been attempted and detected, which often makes it difficult to consult the perpetrator on his cognitive processes leading up to the attack. There is also the problem that successful attacks may never be detected, and therefore the technique for such an attack might never be revealed. This is why material from reformed former practitioners of social engineering, such as Kevin Mitnick, is so valuable. Research can also run into ethical issues, as human subject testing review boards tend to frown on experiments that, as a fundamental part of the experiment, seek to deceive the participants as to the true nature of the study. Also, any research of a phenomenon requires an effective measurement tool, and while one has been proposed for social engineering (Nohlberg, 2005), it is problematic in that it requires the users to be deceived as to the true nature of the study, and that it has yet to be validated through repeated experiments. That problem may not be able to be easily overcome, as deception is central to the success of social engineering, and it is the individual's susceptibility to deception that a researcher would be trying to obtain. Without an effective measurement instrument, it becomes quite difficult for researchers to determine whether a proposed defensive mechanism is effective, because there is no means to compare the before and after. The cumulative effect of all these challenges combined has left the academic literature on social engineering a relatively barren landscape.

1.6 Purpose Statement

Given the relative dearth of research in social engineering when compared to the research on persuasion, research that defines the relationship between social engineering and illegitimate persuasion might allow techniques described by researchers on how to bolster an individual's resistance to illegitimate persuasion to be applied towards increasing an individual's resistance to social engineering. The first purpose of this paper is to establish a link between social engineering and illegitimate persuasion. If this link can be established, it stands to reason that techniques demonstrated in social psychology to be effective against illegitimate persuasion can be incorporated into a training program that can be expected to increase an individual's resistance to social engineering. If the link is shown to be weak or non-existent, then another approach to developing a program to reduce the organization's vulnerability to social engineering attacks will need to be developed.

1.7 Thesis Organization

This paper is organized as follows: Chapter two reviews a variety of relevant research that has been done on social engineering, persuasion and resistance to persuasion. Chapter three details the comparison that was made between social engineering and illegitimate persuasion. Chapter four outlines how techniques designed to reduce an individual's vulnerability to a social engineering attempt can be applied from research on persuasion. Chapter five discusses areas for follow up study as well as potential other areas of application for this research.

II. Literature Review

This chapter is an overview of the concepts associated with social engineering and persuasion. First, the background of social engineering and how psychology previously has been researched in relation to social engineering is reviewed. Persuasion models are then reviewed to give some background on how persuasion might tie into social engineering. Finally, literature that deals specifically on how to develop a resistance to persuasion is reviewed for its potential applicability to social engineering.

2.1 What is meant by the term, "Social Engineering"?

Social engineering has been broadly defined as a set of behaviors used by criminals to surreptitiously gain access to proprietary and confidential information systems (Chantler and Broadhurst, 2006, Granger, 2001). Some behaviors defined as social engineering may be unobtrusive, with no personal interaction occurring between the perpetrator and the information system user. Some examples of unobtrusive behavior that can be labeled as social engineering include dumpster diving and reading the names off a building's information board. While these activities may fit into the broad definition of social engineering, this paper will focus on social engineering behaviors that are more obtrusive, where the perpetrator interacts directly with the victim, using psychological tricks or manipulations to deceive the victim into volunteering information or somehow yielding access to the system. This kind of social engineering can be anything from an authoritative phone call placed to an unsuspecting user, to impersonating tech support and paying a personal visit to "install the latest patches." For the purposes of this paper, the definition of social engineering will refer to this personal approach to social engineering, for two reasons. First, it is the more insidious and pervasive type of social engineering security professionals face today and second, it is the one that might be effectively countered with the appropriate level of training and education.

2.2 Trade Literature and Social Engineering

Much has been written in trade literature about social engineering. Works such as Granger's articles on securityfocus.com (2001, 2002, 2006), Gragg's article published by the SANS Institute (Gragg, 2003), and a book by one of the most infamous hackers that has since reformed, Kevin Mitnick (Mitnick and Simon, 2002), clearly define the threat that organizations face when trying to protect sensitive information. Mitnick, in particular, offers numerous anecdotes to illustrate the various techniques and methods that a social engineer can employ to gather the sought-after information (Mitnick and Simon, 2002). While most trade literature on social engineering does an excellent job of outlining the threat that is faced, the techniques described to combat the threat are fairly universal throughout trade literature.

2.3 Current Preventive Techniques are Inadequate

The universal theme in the trade literature is that social engineering is a problem that is never given enough attention by organizational decision makers (Granger, 2006, Peltier, 2006) and that the best way to defend against it is through user awareness and stringent security policies (Chantler and Broadhurst, 2006, Gragg, 2003, Granger, 2006, Mitnick and Simon, 2002, Peltier, 2006, Rusch, 1999). Authors such as Thornburgh emphasize well-defined policies that establishes an information classification system that clearly defines not only what information can be disclosed but to whom and by whom as well as education policies that instruct users about what information a social engineer can use, how she gains it, and how it can be used towards nefarious ends (Thornburgh, 2004). While these programs and policies certainly are useful at face value in developing a defense against social engineering, the problem is that they have been shown to be inadequate, to the point that many security consultants consider it a given that they will be able to access critical information via social engineering (Gragg, 2003). This is backed up from discussions with individuals of red teams who engage in social engineering behaviors regularly in the course of their attempts to penetrate a network, and who reveal that their attempts to gain information and/or access via social engineering are effective 100 percent of the time. Usually the victims are all too willing to give up the desired information, and the few who actually do challenge the red teams are usually circumvented simply by changing the attack vector, either by attacking another individual, or by switching tactics.

2.4 Psychological Triggers Introduced

Several researchers have tried to better understand why current efforts are inadequate against the threat of social engineering. Some of these efforts have focused on the psychological vulnerabilities of individuals (Chantler and Broadhurst, 2006, Gragg, 2003, Rusch, 1999). Gragg (2003) perhaps made the largest contribution to the literature on the psychological aspects of social engineering when he defined seven psychological triggers that perpetrators can potentially use to manipulate a victim into revealing the desired information. He labeled these triggers as follows: strong affect, overloading, reciprocation, deceptive relationships, diffusion of responsibility, authority, and integrity and consistency. Gaining an understanding of these triggers as well as the techniques used to carry out a social engineering attack is fundamental to being able to develop a defense against the pervasive threat of social engineering.

These triggers are all based on deception, and used by a perpetrator towards a single end: convincing the victim that his or her intentions are benevolent and should be trusted. Once the victim is sufficiently convinced, the perpetrator generally has no problems obtaining the desired information from the victim. Persuasion and deception are central to a successful social engineering attack, and for that reason, understanding persuasion, deception, and the resistance of persuasion is central to finding an effective defense against social engineering. An investigation into social psychology may reveal the best answer as to why previous methods of instilling a resistance to social engineering have proved inadequate and to how best combat the threat of social engineering.

2.5 Cialdini and the Six Principles of Persuasion

When studying the effects and methods of persuasion, researchers over the years have given priority to the study of successful persuasion (getting someone to say yes), rather than on the study of resistance to persuasion (getting someone to say no). This is primarily due to an overwhelming interest and sponsorship by advertisers in getting the individual to say yes to a particular product or service. One of the results of such work on persuasion was Cialdini's (1993, 2001) six persuasion principles that can be used to describe nearly any attempt at persuasion. These principles describe the tendency of individuals to react favorably to stimuli related to each principle. These principles have been described as "heuristic rules" that can assist individuals in their decision whether or not to comply with a request (Cialdini, 1993). The six principles described by Cialdini are authority, consistency, liking, reciprocity, scarcity, and social proof. The following paragraph describes each in more detail.

The principle of authority states that when making a decision, it is common to look to expert advice from an acknowledged source. Authority can be represented effectively simply through the external appearance of authority through specific symbols, such as a uniform, professional title, etc (Bickman, 1974). Another factor that can affect the likelihood of compliance is the legitimacy and credibility of the authority figure. The principle of consistency is based on the desire of individuals to uphold their stated commitments. This tendency is even stronger when a person's values are identified first and the communicator is then able to point out that the request is consistent with those values (Cialdini, 2001). People feel obligated to back up their initial commitment with action, and this tendency can be used as a means of persuasion. The principle of liking refers to the tendency of people to respond positively to a request from someone they like. There are numerous factors that can enhance the effect, such as similarity of attitude, background, physical attractiveness, dress, and the use of praise (Cialdini, 1993). The principle of reciprocity is based on the tendency of individuals to respond in kind when positive behavior is received, such as the giving of gifts, favors, or services. The principle of scarcity refers to the fact that as items or opportunities are perceived to be more scarce, they are perceived to be more valuable, and thus can be leveraged into convincing an individual into a course of action. The last principle that Cialdini describes is the principle of social proof, which refers to the tendency of people to follow the lead of their peers, and that others' standards are often adopted as an individual's own. This can be leveraged into a willingness to comply with a request when supported by the belief that similar peers comply with it as well. These six principles appear to have some significant overlaps with Gragg's psychological triggers. Cialdini's principles and Gragg's triggers are compared in depth and in parallel in chapter four.

2.6 Resistance to Persuasion

Given the focus of this research, our emphasis will be on demonstrated methods for resisting persuasion. Tormala and Petty (2004) define resistance to persuasion as "the process of defending one's attitude against persuasive attack." Other mechanisms through which individuals resist persuasion that have been studied include counter arguing, which Brock described as the "act of generating behavioral cognitive responses against a message" (Brock, 1967). Lydon, Zanna, and Ross (1988) contributed to the resistance to persuasion literature by showing that it was related to bolstering an individual's initial attitudes. It has also been shown that individuals resist persuasion when they have been forewarned (Papageorgis, 1968, Petty and Caciappo, 1977), and when they have been previously exposed to initial attacks of persuasion (McGuire, 1964).

More recently, Tormala and Petty (2002) showed that when individuals resist a persuasive attack, those attacks can actually make the target's ability to resist attack stronger than it already was. These finding were consistent with McGuire's inoculation theory (McGuire, 1964), which showed that exposure to an initial persuasive attack can increase an individual's resistance to subsequent attacks. McGuire's research focused on cultural truisms, which he believed to be a particular vulnerable set of beliefs that people had little experience defending against. He found that individuals who were given the motivation to defend their attitudes, and the ability to do so effectively by a defense that effectively refuted a weakened argument for a cultural truism, were much more likely to have a significant and long-lasting resistance. McGuire's research assumed, though, that initial attitudes were unchanged. Where Tormala and Petty differ is when they showed that those initial attitudes do indeed change, and that the credibility of the source of a persuasive attack affects the predictability of an individual being able to resist a persuasive attack (Tormala and Petty, 2004).

Tormala and Petty demonstrated (2004) that when an attack from what is perceived to be an expert is successfully resisted, an individual's attitude of certainty of being able to resist future attack is increased. When applied towards training individuals and creating a predictable behavior in them, this increased certainty could be useful, leading trainees to a sense of invulnerability. According to Sagarin, et al (2002), the greatest increase in resistance occurs when that sense of invulnerability is shown to be erroneous. Sagarin and his colleagues expanded on McGuire's inoculation theory, and demonstrated in a series of three experiments that instilling resistance requires more than just asserting an individual's vulnerability. Instead, the individual's vulnerability must be demonstrated in order to significantly increase their resistance—that is, the person must have an "a ha!" moment. Applying this technique to social engineering, means that in order to effectively train an individual to be resistant to social engineering, that individual must be shown that he or she is vulnerable to the techniques used by social engineers. Using the techniques developed by Tormala and Petty to ensure that an individual is confident in their ability to resist an expertly executed persuasive attack, and thus seemingly invulnerable to further attacks sets a perfect stage for dispelling that illusion of invulnerability and thus increasing their resistance to further persuasion.

Sagarin and his colleagues conducted three different experiments on resistance to illegitimate persuasion, and the results of each experiment has lessons that can be used towards instilling a resistance to social engineering (Sagarin et al., 2002). I will take a moment here to describe the experiments, the results of each, and the potential application to social engineering.

2.7 Sagarin's Experiments and Their Application to Social Engineering

The first experiment attempted to "instill resistance by teaching participants a rule for discriminating between legitimate and illegitimate appeals and by suggesting to participants that ads containing illegitimate authorities are attempts to deceive consumers" (Sagarin et al., 2002). The results of this experiment suggested that simply giving trainees a discriminating rule did not make them significantly more resistant to an authority-based appeal. The participants did become more discriminating about discerning an illegitimate authority from a legitimate authority, but the experiment rather showed that legitimate authorities came to be seen as more persuasive, rather than instilling a strong resistance against illegitimate authorities (Sagarin et al., 2002). The extension and application for this finding for the purposes of training individuals about social engineering is that teaching individuals about the various psychological triggers simply isn't enough to instill an effective resistance and defense to social engineering. This finding explains why education programs that simply seek to inform the user about the dangers of social engineering have proven ineffective.

The second experiment that Sagarin and his colleagues conducted essentially replicated the first experiment, with one key difference: persistence of the desired effect over time (Sagarin et al., 2002). In addition to testing the participants' resistance to illegitimate persuasion immediately after applying their treatment, the second experiment also tested the participants' resistance one to four days after the treatment was applied. The findings of this second experiment matched the findings of the first experiment in that the participants were found not to have as much an increase in resistance against illegitimate authorities as an increased perception that the legitimate authority was more persuasive. The participants who received the treatment did not resist illegitimate authorities more effectively than did the controls (Sagarin et al., 2002). For the purpose of applying Sagarin's techniques to social engineering, there were two important findings from the second experiment. First, the effects of the treatment did appear to persist over time (Sagarin et al., 2002). That the effects persisted over time suggests that a more formal, interactive training program against social engineering could be effective for a time after being taught. Second, that the participants appeared to recognize the characterization of illegitimacy in the treatment, but appeared to believe themselves unsusceptible to its effects and thus did not act on it (Sagarin et al., 2002). This illusion of invulnerability has been studied in areas from contraception (Burger and Burns, 1988), to mental health (Taylor and Gollwitzer, 1995), to preventative health behaviors (Weinstein, 1987) and has been found to persist even to the individual's detriment. This illusion of invulnerability helps to explain why individuals continue to be vulnerable to social engineering attacks even after the current techniques of training are administered.

The third experiment that Sagarin and his colleagues conducted sought to dispel these illusions of invulnerability by demonstrating to the participants who had been persuaded by an illegitimate authority that they had been fooled (Sagarin et al., 2002). By pointing out to the participants that they had been duped, the participants learned that they were personally susceptible to the risk of illegitimate persuasion. The results of the experiment suggested that participants who had their personal vulnerability demonstrated to them were significantly more resistant to ads containing illegitimate authorities than those who did not have their personal vulnerability demonstrated, and were again more persuaded by legitimate authorities. This experiment was the first time in the series of experiments that the resistance effect was of greater magnitude than the enhancement effect (Sagarin et al., 2002). This demonstrated conclusively that instilling resistance required more than merely asserting an individual's vulnerability; it requires a clear demonstration of this vulnerability (Sagarin et al., 2002).

III. Methodology

This chapter outlines the methodology used to formally compare social engineering and illegitimate persuasion. It provides the necessary information to formalize the relationship in a systematic and definable manner.

3.1 Problem Definition

This research addresses two specific areas of concern about social engineering: determining whether or not social engineering can be equated to illegitimate persuasion, and how to apply techniques shown to be effective against illegitimate persuasion to the problem of training someone to be resistant to social engineering.

3.2 Goals

In order to be able to apply Sagarin's work on illegitimate persuasion to social engineering, the goal of this research is to determine whether or not social engineering is actually comparable enough to illegitimate persuasion to be able to make the transfer of this application to social engineering.

3.3 Approach

This comparison was performed in two phases. First, Gragg's psychological triggers were directly compared to Cialdini's principles of persuasion. This was done to determine whether techniques of persuasion are the same as the techniques of social engineering. Then in the second phase, a systematic, progressive comparison was performed between each of the following aspects of the two: definitions, targets, objectives, and methods of illegitimate persuasion and social engineering, as shown in Figure 2. The experiments that Sagarin and his colleagues performed during their research on illegitimate persuasion are analyzed for the principles of persuasion methods that they used, and for the comparable psychological triggers that were used. One set of experiments focused on the persuasion principle of authority, while the other experiment focused on the principle of scarcity. If the principle of dispelling the illusion of invulnerability works against two separate principles of persuasion, and those principles of persuasion are equated to the psychological triggers of social engineering, then the principle of dispelling the illusion of invulnerability should be applicable to all of the principles of persuasion and, by extension, any psychological triggers that the principles of persuasion are found to be the same as.

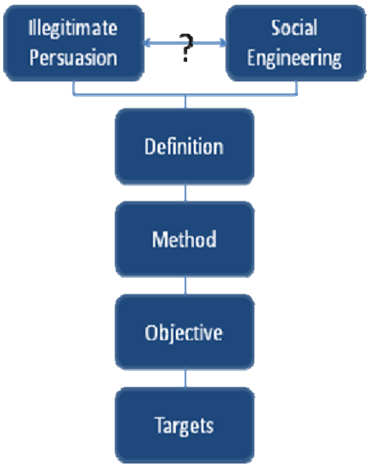

FIGURE 2: COMPARISON METHODOLOGY

This comparison sets the foundation for the training proposal that follows in chapter five and would indicate that the development of a training program designed to instill a resistance to social engineering based on Sagarin's principles of dispelling the illusion of invulnerability in individuals is a viable method of reducing an individual's resistance to social engineering. This type of training would be a revolutionary approach to developing a defense against social engineering, which up to this point has involved comprehensive security policies and education on the methods of techniques and tactics of social engineers, but the current education plans does not address what is a necessary step towards increasing someone’s resistance to a social engineering attempt: the cognitive recognition by the individual that they are vulnerable to the techniques and tactics commonly used by social engineers, and that the information that individual possesses is valuable to a social engineer, and likely to be targeted.

IV. Results and Analysis

4.1 Comparing Psychological Triggers to Principles of Persuasion

In order to compare Gragg’s psychological triggers to Cialdini’s principles of persuasion, each of Cialdini’s principles will be described, followed by any comparable trigger that Gragg describes. This comparison will assist in the analysis of the methods of both illegitimate persuasion and social engineering. Gragg describes seven psychological triggers of social engineering, while Cialdini describes six principles of persuasion, shown in Figure 3.

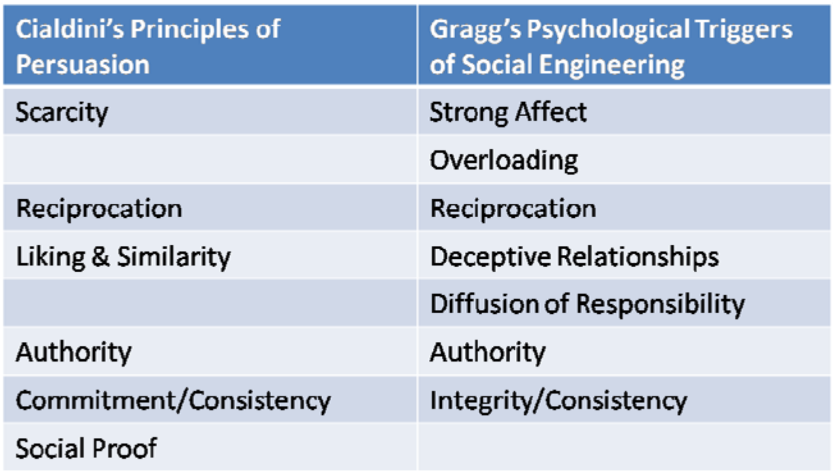

Cialdini’s Principles of Persuasion	Gragg’s Psychological Triggers of Social Engineering
Scarcity	Strong Affect
	Overloading
Reciprocation	Reciprocation
Liking & Similarity	Deceptive Relationships
	Diffusion of Responsibility
Authority	Authority
Commitment/Consistency	Integrity/Consistency
Social Proof	

FIGURE 3: CIALDINI'S PRINCIPLES OF PERSUASION & GRAGG'S PSYCHOLOGICAL TRIGGERS

The first psychological trigger described by Gragg is *strong affect* (Gragg, 2003). This is a technique that uses a heightened emotional state that allows the perpetrator to get information from a victim that would not normally be released. According to Gragg (2003), when employing this technique, the perpetrator will trigger the desired emotion, such as greed, at the outset of the encounter to distract the victim, and interfere with their ability to evaluate, think logically, or develop counterarguments. A good example of a social engineering attack that uses principles of strong affect is the classic phishing attack that claims that a user's banking account information has been compromised. The strong affect emotion this type of attack seeks to elicit is one of fear that their financial information is in the hands of a thief.

The first principle of persuasion that was examined was the principle of *scarcity* (Cialdini, 2001). Scarcity refers to the natural inclination to desire things that are rare or scarce. It can also refer to the desire to miss out on future opportunities. An example of a persuasion attempt that uses the principle of scarcity is an advertisement that touts a "limited-time offer" or "while supplies last." The implication of these examples is that there is a finite amount of time that the product will be available, and any purchase of said product should be done quickly. Like the social engineer in Gragg's trigger of strong affect, an advertiser who uses the scarcity principle to persuade is hoping to elicit a strong emotion that temporarily impairs their judgment. In analyzing the relationship between the trigger of strong affect and the principle of scarcity, it would seem that strong affect can be more broadly applied, and that nearly every circumstance of scarcity fits into the trigger of strong affect.

The second trigger Gragg describes is *overloading*; this is a technique where the perpetrator presents false premises that are laced with convincing truisms to the victim in such rapid succession that the victim becomes mentally passive or when a perpetrator argues from an unexpected perspective and the victim becomes more willing to accept arguments that should have been challenged (Gragg, 2003). There is a time element with the overloading trigger during which the victim is either passive or compliant and the perpetrator acts quickly to take advantage of the victim's reduced motivation to resist. This attack is usually conducted over the phone or, even better, in person to reduce the victim's ability to react and retreat before being forced into making a decision. It can also be used to gain compliance in a small task and set up an attack using a different trigger (such as reciprocation). None of Cialdini's principles of persuasion appear to match this trigger.

The third trigger that Gragg describes is that of *reciprocation*, a technique based on the idea of social exchange where individuals feel obligated to reciprocate a kind gesture from someone else (Gragg, 2003). Also referred to as "reverse social engineering" (Granger, 2001), an example of this type of reciprocation is when the perpetrator creates an IT problem that baffles the user, but he or she just happens to be able to fix. When the perpetrator voluntarily helps to fix the problem, he or she asks the victim to reciprocate by innocently requesting otherwise guarded information.

Reciprocation is a well-established persuasion principle, and is described by Cialdini using the exact same term (Cialdini, 2001). He describes it as a response to positive behavior with positive behavior return, linked to the human need to establish strong social networks (Cialdini, 2001). The caveat with this principle is that according to Brehm's theory of reactance, the initial behavior has to appear genuine, given upfront, and unconditionally or the recipient quickly becomes defensive and guarded (Brehm, 1966).

The fourth trigger that Gragg describes is the *deceptive relationship*, a technique where the perpetrator purposefully builds up a relationship with the intent to eventually extract information out of the victim(s) (Gragg, 2003). Mitnick and Simon (Mitnick and Simon, 2002) describe deceptive relationships with an anecdote where an employee who already was suspicious of Mitnick in a different context was conned. Mitnick took on an alias via email and developed a relationship with the employee by sharing with the employee his distrust and negative impressions of “Kevin Mitnick”. With an established relationship, Mitnick was able to pump the employee for all kinds of information (Mitnick and Simon, 2002).

This trigger fits well with the persuasion principle that Cialdini terms *liking* (Cialdini, 2001). Liking is described as a positive connection between people, and can be fostered in a number of ways including physical attractiveness, similarity, compliments and cooperation (Cialdini, 2001). This principle and the trigger of deceptive relationships are based on the same desired attitudes from the individual, and can be considered one and the same.

The fifth trigger Gragg describes is the *diffusion of responsibility and moral duty*, a technique where the social engineer manipulates the target into believing that he or she will not be held solely responsible for his or her actions and that those actions will have a significant positive consequence, such as saving an employee or helping the company (Gragg, 2003). This trigger is one where the victim realizes the action being solicited is against policy or otherwise prohibited, but proceeds anyway because of a perceived greater benefit. Based on this description, this trigger does not appear to have a corresponding principle of persuasion.

The sixth trigger described by Gragg is response to *authority*. This is a technique that preys on an individual’s propensity to respond and comply with someone who claims to be an authority figure (Gragg, 2003). Compounding the effectiveness of this trigger is that verifying the legitimacy of the authority is considered a challenge, and one that most people are not willing to make. Organizations with very rigid hierarchal structures are particularly vulnerable to this type of attack, the military being the most obvious example. A social engineer might exploit this trigger by pretending to be a high-ranking officer in the military, or some other authority figure and simply request the information he desires. This trigger is another example of one that overlaps perfectly with one of Cialdini's principles of persuasion. Cialdini's principle notes that it is common to seek expert advice, and specific symbols of authorities such as uniforms and professional titles only serve to enhance the effect (Cialdini, 2001).

Finally, the last psychological trigger that Gragg describes is *integrity and consistency,* which takes advantage of the target's desire to follow through on commitments in order to manipulate them into performing some action (Gragg, 2003). This technique targets an individual's perceived obligation to follow through on promises, even if the promises made were not their own. For example, if an employee were to go on vacation and a perpetrator presented a victim with a promise by the employee to provide specific information, the victim will often give up the information if they can be convinced (or convince themselves) the promise was legitimate. This trigger also covers the nature of people to trust others because they believe the person is telling the truth, based on their own tendency towards honesty. This last psychological trigger matches Cialdini's principle of persuasion of *commitment and consistency*, which is linked to an individual's desire to be, or at least appear to be, consistent (Cialdini, 2001). Cialdini notes that this desire is even stronger when a person's values are identified first. For this principle to work well, an initial commitment has to be obtained, and a social engineer could exploit this using reciprocation, or overloading.

4.2 Comparing Definitions of Illegitimate Persuasion and Social Engineering

The definitions used for comparing illegitimate persuasion and social engineering were taken from two of the most influential articles about each subject. Dr. Sagarin's and his colleagues' definition of illegitimate persuasion was used in order to be consistent, as his article was the basis of this comparison. Sagarin defines illegitimate persuasion as "messages that seek to convince through deceptive tactics” (Sagarin et al., 2002). In the article that Gragg describes the psychological triggers of social engineering, he defines social engineering as “the process of deceiving people into giving away access or confidential information” (Gragg, 2003).

When comparing Gragg’s definition of social engineering and Sagarin’s definition of illegitimate persuasion, what immediately jumps out from these two definitions is that deception is fundamental to both. Sagarin’s definition of illegitimate persuasion implies that the deception is purposeful and intentional, while Gragg’s definition of social engineering leaves less room for interpretation. By his definition, social engineering is a designed, purposeful “process” that leaves no doubt about whether the social engineer has a malicious intent and design to achieving his goal. Illegitimate persuasion’s definition, however, does not define a malicious intent to this degree. While social engineering’s definition definitely implies a malicious intent and illegitimate persuasion’s definition does not, the fundamental premise to both is deception. When developing a training program against social engineering, the goal of that program is to get the user/participant to recognize when they are being deceived, so finding that deception is fundamental to illegitimate persuasion as well is positive, and allows the next comparison to be made.

4.3 Comparing Objectives and Targets

Next, an examination was made on the relationship between the objectives and targets of illegitimate persuasion and social engineering. In all three of their experiments, Sagarin and his colleagues had their participants examine full-page color advertisements from periodical magazines with authority figures peddling their products (Sagarin et al., 2002). The objective of illegitimate persuasion in the context in which Sagarin describes it (advertising) is to convince the individual who receives the message to purchase a particular product or service. The distinction between a legitimate and illegitimate advertisement was made based on whether the individual recommending the product had reasonable credentials to do so. For example, a model dressed as a stockbroker recommending the Wall Street Journal was labeled as an illegitimate authority, as he was not an expert on the product he was trying to sell. Alternately, a legitimate authority was seen as someone who did have expert credentials for the product they were selling, such as Marcel Cockaerts, President of Kredietbank, Brussels recommending insurance for banks (Sagarin et al., 2002). Social engineering's objective by definition is to convince an individual to voluntarily give away some sort of access or critical information. The objectives of social engineering and illegitimate persuasion clearly cannot be equated on this basis, however breaking the objectives down to their core may allow for a more direct comparison. Illegitimate persuasion and social engineering both seek to compel behavior that the individual likely would not have taken otherwise. The desired behavior of the recipient is different, but ultimately the objective of both is to persuade and individual to take an action, and in that way the objectives are very much aligned. So, while the objectives of illegitimate persuasion and social engineering are not exactly alike, they do appear to be similar enough that it still stands to reason that research shown to be effective in bolstering a resistance to illegitimate persuasion can be applied to social engineering. Once the objectives were compared, the next step was to examine the targets.

The target of an illegitimate persuasion attempt is the advertising audience at which the marketing campaign is directed. In the experiments that Sagarin and his colleagues ran, the targets that were chosen to be subject to illegitimate and legitimate authority advertisements included the demographic of undergraduate college students (Sagarin et al., 2002).

Advertisements are often targeted to reach a very specific demographic, and the advertisements reflect the values and desires of the targeted demographic. The target of a social engineering attack is an individual with specific information (or access to that information) that the social engineer desires. A social engineer may target a specific individual, or may "shotgun" an attack out to a broad audience. The latter example describes when a social engineer's target is most like the target of an advertisement that uses illegitimate persuasion. However, because the social engineer's only target may be a specific individual and the goal of an advertisement is to reach as broad an audience as possible within a given demographic, the targets of social engineering and illegitimate persuasion cannot be completely equated. Indeed, the targets of the persuasion attempt may be the single biggest difference between illegitimate persuasion and social engineering. For the purposes of applying the techniques of bolstering resistance to illegitimate persuasion to social engineering, this difference is acceptable because the techniques are not exclusive to a particular demographic or class of individuals. At this point, a comparison has been made between the definitions, objectives, and targets of illegitimate persuasion and social engineering. Next, comparison was made between the methods of the two.

4.4 Comparing Methods & Techniques

Sagarin's research on bolstering resistance focused on the methods used to illegitimately persuade, so in order to be able to apply his methods of bolstering resistance to social engineering, the methods used to accomplish both were compared. The psychological aspects of two specific examples of illegitimate persuasion were evaluated to determine if they fall within one of the known psychological triggers of social engineering (Gragg, 2003). If the techniques used in these illegitimate persuasion examples use Cialdini's principles of persuasion that have been equated to Gragg's psychological triggers from social engineering, then the techniques that have been shown to be effective in bolstering an individual's resistance to illegitimate persuasion should be effective against those same principles when the desire is to prevent social engineering.

4.4.1 Authority Techniques

When describing the techniques on how to bolster an individual's resistance to illegitimate persuasion, Sagarin and his colleagues used two types of situations to dispel the individual's illusion of invulnerability (Sagarin et al., 2002). In the first example, the first of a series of three experiments gave the participants rules about determining the legitimacy of authority-based appeals and investigated whether or not that led to an increase in resistance against illegitimate appeals. The first experiment measured the participants' resistance immediately after being shown the difference between illegitimate and legitimate advertisements, while the second experiment repeated the first experiment and measured the participant's responses after one to four days. The third experiment differed from the previous two in that, when a participant incorrectly described the illegitimate authority as legitimate, the error was pointed out to the participant. What all three of these experiments had in common was that each asked the participants to differentiate between a legitimate authority figure, such as a doctor recommending a pain killer; and an illegitimate authority figure, such as an actor dressed as a stock broker recommending a particular trading firm. This persuasion appeal clearly uses Cialdini's principle of authority as it appeals to the idea that the individual recommending the product has a clear, demonstrated expertise in the domain of the advertisement. Similarly, social engineers use an appeal to authority, which has already been linked to Cialdini's persuasion principle of authority through Gragg's trigger of authority, to manipulate a victim into giving up sensitive information. A classic example of a social engineer using authority is when they claim to be technical support, either over the phone or in person. Because the dispelling of illusion of invulnerability techniques were shown to be most effective at instilling a resistance against an authority-based illegitimate persuasion attempt and because the persuasion principle of authority and the social engineering trigger of authority have been shown to be one and the same, those same dispelling of illusion techniques should also be effective at instilling a resistance against the social engineering trigger of authority.

4.4.2 Scarcity Techniques

In a different study on illegitimate persuasion, the principle of persuasion that the study focused on building a resistance against was the principle of scarcity (Coutinho and Sagarin, 2007). In this experiment, Coutinho and Sagarin examined how participants responded to legitimate and illegitimate claims of scarcity, and determined whether the principles of dispelling the illusion of invulnerability would be as effective as in the experiment done on authority. The experiment gave the participants the rule that in order for there to be a legitimate claim of scarcity, it must be apparent that the object could not be produced anymore (such as coins from the 1800s). The experiment used Cialdini's persuasion principle of scarcity to elicit an emotional response from the target that created a sense of urgency to purchase the item. For the experiment, participants were randomly assigned into four groups and each participant was given a series of six advertisements, balanced with three illegitimate scarcity claims and three legitimate scarcity claims. The four groups were then assigned one of the four following conditions: the *control group* which received no training; the *asserted vulnerability* group, which received guidance on how to distinguish between a legitimate scarcity claim and an illegitimate scarcity claim, but received no demonstration of their personal vulnerability, and were not surveyed on how convincing they found the advertisement; the *asserted vulnerability with inquiry condition* group which was the same as the asserted vulnerability group, with the exception they were surveyed on how convincing they found each advertisement; and the *demonstrated vulnerability treatment* group, which were surveyed on how convincing they found the advertisement, and were then told they had been fooled by any illegitimate advertisement they found convincing and given training on how to distinguish between an illegitimate scarcity claim and a legitimate one. A summary of the assigned group and the conditions given to each is shown in Figure 4.

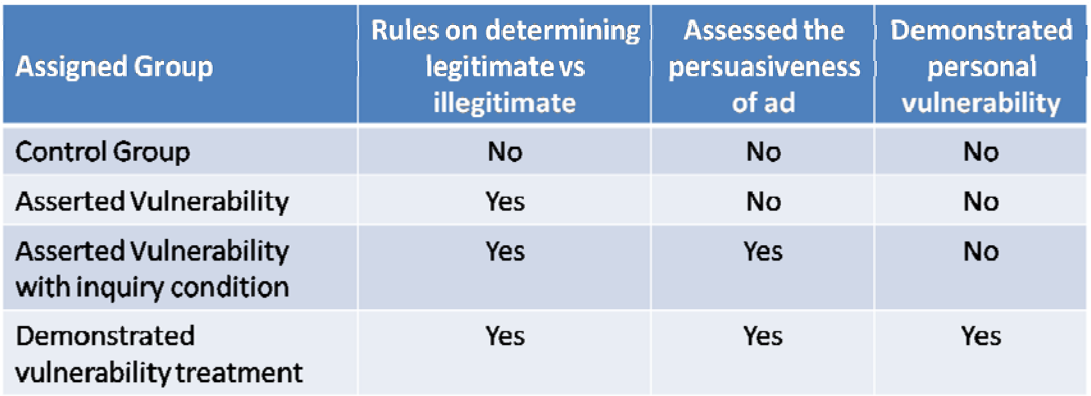

Assigned Group	Rules on determining legitimate vs illegitimate	Assessed the persuasiveness of ad	Demonstrated personal vulnerability
Control Group	No	No	No
Asserted Vulnerability	Yes	No	No
Asserted Vulnerability with inquiry condition	Yes	Yes	No
Demonstrated vulnerability treatment	Yes	Yes	Yes

FIGURE 4: SCARCITY EXPERIMENT GROUPS AND TREATMENTS

After reading the training material, participants then rated six more advertisements to assess on how persuasive the advertisement was and how much the participant perceived an undo manipulative intent. Again, these advertisements were balanced with three legitimate scarcity claims and three illegitimate scarcity claims. The result of the experiment found that the participants of the demonstrated vulnerability treatment group were far more resistant to scarcity claims than the participants of the other groups. This confirms the results of the earlier experiments using the principle of authority. As shown above, the scarcity principle of persuasion can be equated to the strong affect trigger of social engineering, and so the results of this experiment should be able to be extended towards instilling a resistance to social engineering. Furthermore, the result from this experiment suggests that the technique of dispelling an individual’s illusion of invulnerability is necessary to instill a resistance to any of the principles of persuasion, and by extension, the triggers of social engineering.

4.5 The Conclusion: Illegitimate Persuasion = Social Engineering

While illegitimate persuasion and social engineering cannot be exactly equated, the principles of persuasion are very similar to the psychological triggers of social engineering, and it reasonably follows that techniques that have been shown to be effective at reducing an individual’s resistance to illegitimate persuasion will also be effective in reducing an individual’s vulnerability to social engineering. By extension, it follows that a combination of comprehensive security policies, combined with training that incorporates into its curriculum a dispelling of an individual’s illusion of invulnerability, is fundamental and necessary to protecting an organization from being exploited by social engineering attacks. The application of this conclusion to social engineering represents a significant paradigm shift on how information security specialists approach the task of training users to resist social engineering attacks. Rather than simply educating users on the dangers and methods of social engineers, the individual needs to be exposed to training that demonstrates that user’s personal vulnerability in order for the training to be effective. The question then becomes: How do you design a training program that is effective in demonstrating to the user their personal vulnerability to a social engineering attack?

V. Discussion

The previous chapter focused on determining whether or not illegitimate persuasion could be equated to social engineering. This chapter examines the principles necessary to apply these principles to a training program that instills a resistance to social engineering. The Constructivist Learning Theory is reviewed as it was a suggested method by Coutinho and Sagarin as a potential training technique that would be well suited to the task of incorporating the principle of dispelling the illusion of invulnerability into a training program (Coutinho and Sagarin, 2007). Other potential training techniques are also reviewed for their potential applicability. Then the requisite skill level described by Bloom's Taxonomy that an individual would need to be trained up to is examined. Finally, some suggested methods that incorporate training techniques with the requisite skill level, and that apply the principle of dispelling the illusion of invulnerability training are given.

5.1 Constructivist Theory and Its Applicability

Coutinho and Sagarin argued that a training program built on the Constructivist Learning Theory was the best way of building a training program built on the technique of dispelling an individual's illusion of invulnerability (Coutinho and Sagarin, 2007). This approach does seem logical as the Constructivist Learning Theory is built on the premise that students learn through experience and that each new experience builds on their own existing knowledge learned from experience. A training program based on this learning theory would use the dispelling the illusion of invulnerability technique to construct a new experience or set of experiences that would build on the trainee's own set of experiences by demonstrating that individual's specific vulnerability. Training based on the Constructivist Learning Theory would incorporate examples of social engineering into scenarios based on each of the seven different psychological triggers, and then build on those triggers by demonstrating common ways the triggers can be used in combination. An example of such a combination might be using the strong affect trigger to get a small commitment from an individual to give up an innocuous bit of information, and then exploiting that commitment by using the trigger of integrity and consistency to gain access to information that is more protected. A training program that incorporates scenarios such as this into its curriculum would give the individual receiving the training the opportunity to construct a new set of experiences that he or she would retain and apply in their day-to-day life and job.

5.2 Other Training Techniques

While the principle of dispelling the illusion of invulnerability may be best implemented by techniques build on the Constructivist Learning Theory, it would be impractical to build a scenario for each possible technique that an individual might be approached by someone attempting to perpetrate a social engineering attack. Such training should serve as a technique to make the individual more receptive to an education that also incorporates pattern recognition and the ability to evaluate a situation for its potential for a social engineering attack to take place. Current methods of informing individuals on the threat of social engineering as well as the techniques used by social engineers certainly have their place in a training program designed to reduce an individual's vulnerability to social engineering.

5.3 Bloom's Taxonomy, Its Revision, and Its Applicability

It should also be considered as to what level of skill is necessary for an individual to gain as a result of going through a training program designed to instill a resistance to social engineering. Bloom's Taxonomy is a universally recognized framework for classifying what is expected of students as the result of instruction. Bloom's Taxonomy describes the following six levels of competency, in order from lowest to highest complexity: knowledge, comprehension, application, analysis, synthesis, and evaluate. In 2001, Krathwohl introduced a revised version of Bloom's Taxonomy that changed the terms into gerunds to better describe the expected ability of the student at each level, and swapped the order of synthesis and evaluate, as synthesis is widely regarded as the highest level of competency an individual can achieve (Krathwohl, 2002). Krathwohl's revision of Bloom's Taxonomy describes the following six levels of competency in order of complexity, from lowest to highest: remember, understand, apply, analyze, evaluate, and create. A comparison of the original taxonomy and Krathwohl's revision can be seen in Figure 5.

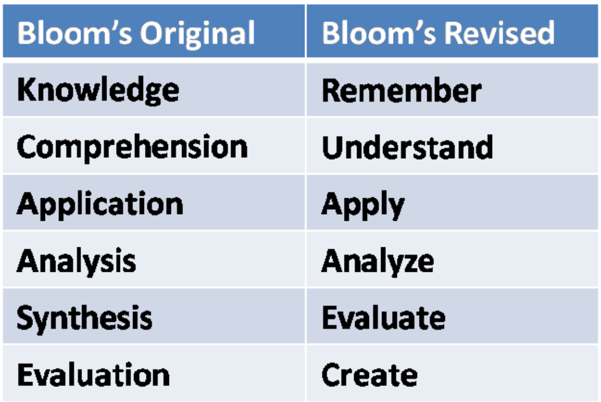

Bloom's Original	Bloom's Revised
Knowledge	Remember
Comprehension	Understand
Application	Apply
Analysis	Analyze
Synthesis	Evaluate
Evaluation	Create

FIGURE 5: BLOOM'S ORIGINAL AND REVISED TAXONOMY (LOWEST TO HIGHEST)

Currently, the typical educational program for preventing a social engineering attack expects its students to leave simply with an increased awareness about the techniques and tactics that social engineers employ. In the revised taxonomy, the level of competency that such programs teach to is the understand level. While this initially seems sufficient, the implications of Sagarin's and Coutinho's research described above are that rule-giving is not enough, and that it must be demonstrated to the individual that they are personally vulnerable. This demonstration actually educates the user to a higher level of competency, the evaluate competency. Once their vulnerability has been demonstrated for them, the individual can then make judgments based on established criteria and evaluate each new situation that is presented to them. Therefore, the target level of competency for training that seeks to instill a resistance against social engineering should be the evaluate competency.

5.4 What should be in a Training Program

Security policies must be able to identify critical information that needs to be protected. Why the information is being protected must make sense to the user, and then formal rules must be put in place that forbids the user to release the identified critical information. There must be "buy-in" from the user that there is an urgency to protect the information, and this can be achieved through training. A potential training example might be to give the user a series of questions to determine whether or not specific information was ok to release. The sample set should include obvious examples, such as username and password, and then innocuous examples such as the individual's name and rank.

User/employee training should include samples of behavior and information that might appear to be innocuous, but in fact are not. When they incorrectly choose an action or information that has been identified as critical, the training should immediately demonstrate or present an historical example of how that that type of activity was used to perform malicious activities. This type of training uses the technique of dispelling a user's illusion of invulnerability to create a salient, *teachable moment* that can affect cognitive processes and ultimately their influence behavior so that they increase their vigilance.

Critical to this training is that employees are demonstrated two things: First, everyone, regardless of their position within the organization, is vulnerable to a range of social engineering techniques. Secondly, each person must be made aware that the social engineer is targeting people just like him or her. The salience of the demonstration as well as evidence of the individual's vulnerability will bolster that individual's resistance to social engineering attacks in the future. In order to effectively create a resistance to a variety of attacks, this technique should be demonstrated to the user across all the psychological triggers of social engineering.

5.5 Limitations of Applicability & Potential for Future Research

The finding of this thesis, and the effectiveness of any training program created to incorporate this research depends largely on the ability of previous research done on dispelling the illusion of invulnerability against two different principles of persuasion. Applying the principle of dispelling the illusion of invulnerability to principles of persuasion other than authority and scarcity has not been performed, and thus the potential exists that the techniques may not be effective against six principles, much less all seven psychological triggers. With that said, more research is definitely required in this area.

If a means of measuring an individual's vulnerability to social engineering can be developed, then a study that verifies the effectiveness of this type of training when applied to social engineering would add validity to this study and to the transferability of Sagarin's work to the problem of countering social engineering attacks. Perhaps such research could incorporate methods of measuring resistance from Sagarin's work and apply them to social engineering. Research that measures an individual's resistance prior to treatment, gives them a treatment, and then measures their resistance post treatment would help validate the concepts from this study. A longitudinal study that determines the rate of decay of training, and how long the treatment is effective for would also be useful for determining how often individuals have to be subjected to training.

Bibliography

BICKMAN, L. (1974) The Social Power of a Uniform. *Journal of Applied Social Psychology,* 4**,** 47-61.

BREHM, J. W. (1966) *A Theory of Psychological Reactance,* New York, Academic.

BROCK, T. C. (1967) Communication discrepancy and intent to persuade as determinants of counterargument production. *Journal of Experimental Social Psychology,* 3**,** 296-309.

BURGER, J. M. & BURNS, L. (1988) The Illusion of Unique Invulnerability and the Use of Effective Contraception. *Personality & Social Psychology Bulletin,* 14**,** 264-170.

CHANTLER, A. & BROADHURST, R. (2006) Social Engineering and Crime Prevention in Cyberspace. Queensland University of Technology.

CIALDINI, R. B. (1993) *Influence: Science and practice,* New York, Harper Collins.

CIALDINI, R. B. (2001) Harnessing the science of persuasion. *Harvard Business Review.*

COUTINHO, S. & SAGARIN, B. J. (2007) Instilling Resistance to Scarcity Advertisement. *Studies in Learning, Evaluation, Innovation and Development,* 4**,** 54-66.

GRAGG, D. (2003) A Multi-level Defense Against Social Engineering. SANS Institute - as part of Information Security Reading Room.

GRANGER, S. (2001) Social Engineering Fundamentals, Part I: Hacker Tactics. *Infocus.*

GRANGER, S. (2002) Social Engineering Fundamentals, Part II: Combat Strategies. *Infocus.* Security Focus.

GRANGER, S. (2006) Social Engineering Reloaded. *Infocus.* Security Focus.

KRATHWOHL, D. R. (2002) A Revision of Bloom's Taxonomy: An Overview. *Theory Into Practice,* 41**,** 7.

LYDON, J., ZANNA, M. P. & ROSS, M. (1988) Bolstering attitudes by autobiographical recall: Attitude persistence and selective memory. *Personality and Social Psychology Bulletin,* 14**,** 78-86.

MCGUIRE, W. J. (1964) Inducing resistance to persuasion: Some contemporary approaches. *Advances in Experimental Social Psychology,* 1**,** 191-229.

MITNICK, K. & SIMON, W. (2002) *The Art of Deception,* Indianapolis, Wiley Publishing, Inc.

NOHLBERG, M. (2005) Social Engineering Audits Using Anonymous Surveys - Conning the Users in Order to Know if They Can Be Conned. *4th Security Conference.* Las Vegas.

PAPAGEORGIS, D. (1968) Warning and Persuasion. *Psychological Bulletin,* 70**,** 270-282.

PELTIER, T. R. (2006) Social Engineering: Concepts and Solutions. *Information Systems Security,* 15**,** 13-21.

PETTY, R. E. & CACIAPPO, J. (1977) Forewarning, cognitive responding and resistance to persuasion. *Journal of Personality and Social Psychology,* 35**,** 645-655.

RUSCH, J. J. (1999) The "Social Engineering" of Internet Fraud. *Internet Networking 1999.* San Jose.

SAGARIN, B. J., CIALDINI, R. B., RICE, W. E. & SERMA, S. B. (2002) Dispelling the Illusion of Invulnerability: The Motivations and Mechanisms of Resistance to Persuasion. *Journal of Personality and Social Psychology,* 83**,** 526-541.

TAYLOR, S. E. & GOLLWITZER, P. M. (1995) Effects of Mindset on Positive Illusions. *Journal of Personality and Social Psychology,* 69**,** 213-226.

THORNBURGH, T. (2004) Social Engineering: the "Dark Art". *Proceedings of the 1st annual conference on Information security curriculum development.* Kennesaw, Georgia.

TORMALA, Z. L. & PETTY, R. E. (2002) What Doesn't Kill Me Makes Me Stronger: The Effects of Resisting Persuasion on Attitude Certainty. *Journal of Personality and Social Psychology,* 83**,** 1298-1313.

TORMALA, Z. L. & PETTY, R. E. (2004) Resistance to Persuasion and Attitude Certainty: The Moderating Role of Elaboration. *Personality & Social Psychology Bulletin,* 30**,** 1446-1457.

WEINSTEIN, N. D. (1987) Unrealistic Optimism about Susceptibility to Health Problems. *Journal of Behavioral Medicine,* 10**,** 481-500.

WINKLER, I. S. & DEALY, B. (1995) Information Security Technology?...Don't Rely On It - A Case Study in Social Engineering. *Proceedings from the Fifth USENIX UNIX Security Symposium.* Salt Lake City, Utah.

Vita

Jamison W. Scheeres graduated from Northridge High School in Layton, Utah in 1997. He entered undergraduate studies at the United States Air Force Academy in Colorado Springs, Colorado where he graduated with a Bachelor of Science degree in Social Sciences in May 2001. Upon graduation, he was commissioned as a Second Lieutenant, and assigned to the 81st Communications Squadron at Keesler AFB, Mississippi where he served in a variety of roles as a communications officer, including creating weekly information security intelligence briefings for 18 months post 9/11, serving as the base software license manager, and serving as supervisor for technical solutions. From Keesler AFB, he was assigned to the Information Technology Wing of the NATO E-3 AWACS Component in Geilenkirchen, Germany. There he served as the supervisor for the configuration management section, as well as the deputy branch chief for the systems and servers branch. In August 2006, he entered the Graduate School of Engineering and Management, Air Force Institute of Technology. He separated from active duty Air Force in January 2008 and upon graduation will work as an Assistant Vice President for Bank of America in Charlotte, North Carolina.

REPORT DOCUMENTATION PAGE	*Form Approved* *OMB No. 074-0188*

1. REPORT DATE *(DD-MM-YYYY)* 03-27-08	**2. REPORT TYPE** Master's Thesis	**3. DATES COVERED** *(From – To)* August 2006 – March 2008

4. TITLE AND SUBTITLE

Establishing the Human Firewall: Reducing an Individual's Vulnerability to Social Engineering Attacks

5a. CONTRACT NUMBER

5b. GRANT NUMBER
08-142

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)

Scheeres, Jamison W.

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAMES(S) AND ADDRESS(S)
Air Force Institute of Technology
Graduate School of Engineering and Management (AFIT/EN)
2950 Hobson Way, Building 640
WPAFB OH 45433-8865

8. PERFORMING ORGANIZATION REPORT NUMBER

AFIT/GIR/ENG/08-04

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
2LT Julie Ann Janson, AFRL/RHX
2255 H Street Bldg 248 WPAFB, OH 45433
Comm: (937) 656-6542
Julie.Janson@wpafb.af.mil

10. SPONSOR/MONITOR'S ACRONYM(S)

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

13. SUPPLEMENTARY NOTES

14. ABSTRACT

Hackers frequently use social engineering attacks to gain a foothold into a target network. This type of attack is a tremendous challenge to defend against, as the weakness lies in the human users, not in the technology. Thus far, methods for dealing with this threat have included establishing better security policies and educating users on the threat that exists. Existing techniques aren't working as evidenced by the fact that auditing agencies consider it a given that will be able to gain access via social engineering. The purpose of this research is to propose a better method of reducing an individual's vulnerability to social engineering attacks.

15. SUBJECT TERMS
Social Engineering; Resistance to Persuasion

16. SECURITY CLASSIFICATION OF:			**17. LIMITATION OF ABSTRACT**	**18. NUMBER OF PAGES**	**19a. NAME OF RESPONSIBLE PERSON** Mills, Robert F.
a. REPORT	**b. ABSTRACT**	**c. THIS PAGE**			**19b. TELEPHONE NUMBER** *(Include area code)* (937) 255-6565, ext 4527 (Robert.mills@afit.edu)
U	U	U	UU	49	